

DARKTRACE

# The AI Arsenal

---

Understanding the tools  
that shape cybersecurity

■ LLM

■ Anomaly detection

■ Graph theory

■ Bayesian meta-classifier

■ Decision tree model

■ Chat

---

# Abstract

Darktrace is a research and development-led company that applies unsupervised and supervised machine learning to the challenge of detecting threatening digital activity within corporate networks, with the aim of combatting novel cyber-attacks that bypass rule-based security tools.

This white paper examines the multiple layers of machine learning that make up Darktrace's Cyber AI and how they are architected together to create an autonomous system that self-updates, responding to – but not requiring – human input.

---

Thanks to

**Tim Bazalgette,**  
PhD, SVP AI / ML, R&D

**Nicole Carignan,**  
SVP, Security & AI Strategy, Field CISO

**Hanah-Marie Darley,**  
MSc, Director, Security & AI Strategy, Field CISO

**Jack Stockdale,** CTO

**Brittany Woodsmall,**  
MBA, Manager, Product Marketing, AI

and **Lily Steinberg,**  
Manager, Content Marketing

for their contributions.

# Table of Contents

03	A new age of cyber-threat
04	Traditional approaches to cybersecurity
04	Strengths and challenges of modern AI techniques
	Supervised machine learning and cybersecurity
	Large Language Models (LLMs) and cybersecurity
	Unsupervised machine learning and cybersecurity
06	Darktrace machine learning: Combining multiple techniques
07	Multi-layered AI
	1. Behavioral prediction
	2. Real-time threat detection
	3. Customizable model editor
	4. Investigation and response
16	Pulling the AI layers together

# A new age of cyber-threat

The last decade has seen an unmistakable escalation in cyber conflict, as criminals, nation states, and lone opportunists take advantage of digitization and connectivity to compromise systems and gain advantage – whether financial, reputational, or strategic.

Cyber threat actors are constantly innovating too, launching new attack tools and technologies to get through traditional defenses, such as firewalls and signature-based gateways, and into the file shares or accounts that are most valuable to them.

**Most recently, with the rise of generative AI, attackers have adopted automation to further increase the volume, velocity, and variety of attacks.**

78%

of CISOs now agree AI-powered cyber-threats are already having a significant impact on their organization.

95%

believe AI can improve the speed and efficiency of cyber defense.

Adversarial use of AI allows for more rapid iterations on advanced Tactics, Techniques, and Procedures (TTPs) to exploit vulnerabilities and evade detection. Also, increasingly popular subscription-based tools such as Malware-as-a-Service (MaaS) and Ransomware-as-a-Service (RaaS) have lowered the barrier-to-entry for less experienced attackers, making it easier to carry out complex, multistage attacks.

Attackers have further exploited the digital complexity of the average organization of today, which shares data across multiple locations, devices, and technology services– from cloud services and Software-as-a-Service (SaaS) tools to untrusted home networks and unofficial Internet of Things (IoT) devices. Meanwhile, malicious insiders remain a constant threat.

---

## **The emergence of these new threats necessitates a technical change:**

How can we use machine learning to detect what we do not know to anticipate? In other words, without relying on data sets of previous attacks, how can we build a system that learns what a threat is, by comparing its behavior to everything else going on in that environment?

This white paper explains Darktrace's approach to AI and machine learning and shines a light on the unique interplay between unsupervised machine learning and supervised machine learning behind the core of our multi-layered AI architecture.



# Traditional approaches to cybersecurity

According to the traditional security paradigm, firewalls, endpoint security solutions, and other tools such as Security Information and Event Management (SIEMs) and Intrusion Detection Systems (IDS) are deployed to enforce specific policies and provide protection against known threats.

While these tools have a part to play in an organization's overall defense posture, they are ill-equipped to tackle the new age of rapidly evolving cyber-threats, especially as enterprise infrastructures diversify.

## Strengths and challenges of modern AI techniques

**Perimeter controls** are narrow – if they miss an attack at the point of entry, they have failed and cannot take further action.

**Edge security** primarily depends on known rules and signatures, specializing in detecting attacks that have been previously identified. While these tools can sometimes detect unknown attacks, it is not their strong suit, making them less effective against novel threats. Another concern with this dependency is the rise of disclosed vulnerabilities of edge security tools observed in 2024 and its expected continuation into the future.

**Log tools and SIEM databases** require manual effort to ensure data is consistently collected across the entire organization and matched against the security team's predictions of threats. They rely on the security team imagining everything that might possibly go wrong, without overwhelming analysts with alarms.

What is typically considered **behavioral analytics** often relies on the rules-based paradigm of configuring how certain job titles or devices 'should' behave and then looking for deviations in that behavior. Or, products use human-led supervised machine learning models to train for suspicious or anomalous behavior, limiting detection to prior, static knowledge. These approaches fail to scale to the complexity and size of modern businesses.

**Ultimately, legacy systems have been outpaced by modern business complexity and attacker innovation, suffering from these fundamental constraints.**

---

### They need to:

- Know about all previous attacks.
- Perfectly understand your business and business-specific rules.
- Have a flawless way of sharing high-quality information about new attacks.
- Predict the threat landscape to achieve zero-trust initiatives.
- Be able to turn all the above insights into rules, signatures, or workflows that work.

Machine learning has presented a significant opportunity to the cybersecurity industry, and many vendors have been using it for years. Although, despite the high potential benefit of applying machine learning to cybersecurity, not every AI tool or machine learning model is the same.

For example, even legacy tools that use AI to add automation to processes still require victims and data from compromises before they can provide solutions. The age of unpredictable, fast-moving attacks has rendered this approach woefully deficient.

Learning about the different types of AI and their roles in cybersecurity leads to a better understanding of what is required for a robust, multi-layered AI security solution. Below is a description of some of the main types.

# Supervised machine learning and cybersecurity

Some cybersecurity vendors have been experimenting with supervised machine learning for over 10 years. Most use big data science, shared cyber-threat intelligence, known or reported attack behavior, and classifiers to be able to automate threat detection-based reported data.

## **Supervised machine models are trained on labeled data.**

In cybersecurity, this often means a database of previously seen behaviors, where each behavior is known to be either malicious or benign and is labeled as such. New activities are then analyzed to see whether they more closely match those in the malicious class or those in the benign class. Any that are evaluated as being sufficiently likely to be malicious are flagged as threats.

## **Supervised machine learning systems are best equipped to give you an explicit answer based on prior knowledge.**

For example, we can feed a system with lots of examples of known ransomware, and it will learn the common indicators of that malware and be able to detect similar attacks in the future. This is an advantage over signature-based approaches, as it is less specific to a given signature and can learn to identify threats based on a mixture of more complex properties.

However, “overfitting” is a common problem in supervised machine learning, where model parameters are too finely tuned to the training data. Instead of learning the essence of a category, the machine learns a particular example – for example, a supervised machine learning classifier may learn to recognize the explicit TTPs of a known ransomware strain, but as adversaries vary their TTPs, use living-off-the-land techniques, or tools that do not match signatures, they will remain undetected by this group. This can lead to these supervised machine learning models missing patterns or features outside of the data set.

In the last five years, we have seen more vendors expand into the behavior analytics and anomaly detection side, but it is limited to a baseline (or static) normality understanding and one or two types of machine learning models to detect anomalies on top of that.

This method is based on predefined behaviors of each device or user within the digital estate. The method separates the learning, when the behavioral profile is created, from the subsequent anomaly detection. As such, it does not learn continuously and requires periodic updating and re-training to try to stay up to date with dynamic business operations and so opens the door for a high rate of daily false positives and false negatives.

**Learn more about the difference between leading AI cybersecurity methods and Darktrace’s unique approach in the journal article “[Anomaly-Based Threat Detection: Behavioral Fingerprinting Versus Self-Learning AI.](#)”**

## **Systems that rely entirely on supervised machine learning have fundamental constraints:**

- If a model is trained to detect a specific kind of activity, then it will be unable to identify new kinds of activity that are sufficiently distinct.
- The training data needs to be labeled, which may require a large amount of human input or other methods like machine labeling or determining labels based on an existing property of the dataset.
- Any mislabeled data or human bias introduced can seriously compromise the ability of the system to correctly classify new activities.
- Supervised machine learning is static by nature and requires regular re-training and updating.

# Unsupervised machine learning and cybersecurity

Another AI method applied to cybersecurity is unsupervised machine learning. Unlike supervised approaches, it does not require labeled training data. Instead, it can identify key patterns and trends in the data, without the need for human input.

Unsupervised machine learning has the potential to independently classify data and detect compelling patterns, instead of relying on knowledge of past threats. This removes the dependency of human input or involvement to guide learning. Accuracy and efficacy will be largely dependent on the data and machine learning technique employed for the specific use case. Depending on the techniques employed, unsupervised machine learning can be used to continuously learn, update, and self-mediate.

As an example, we use unsupervised machine learning techniques to form an understanding of ‘normal’ behaviors across the infrastructure, pertaining to devices, users, or cloud containers and sensors, and it detects deviations from this evolving ‘pattern of life’ that may point to a developing threat.

However, like any machine learning technique, unsupervised machine learning has some weaknesses. It is constrained by input parameters, requiring a thoughtful set up to ensure the accuracy of the outputs. Additionally, while it can discover patterns in data, some of those patterns may be irrelevant and distracting. Finally, the outputs come in the form of anomalies rather than threats, requiring more human interpretation or additional modeling to translate them into insights for the security team.

---

## Large Language Models (LLMs) and cybersecurity

With the recent explosion of LLMs in the market, many vendors are rushing to add generative AI to their products, using it for chatbots, Retrieval-Augmented Generation (RAG) systems, agents, and embeddings. LLMs power generative AI, and they can be used in both supervised and unsupervised ways. They are pre-trained on massive volumes of data and can be applied to human language, machine language, and more.

**In security, generative AI can optimize data retrieval for defenders by converting natural language queries to power queries across tools and different data sets.**

It can also be used to summarize as part of the reporting process or to emulate sophisticated phishing attacks for preventative security. But, since this is semantic analysis, LLMs can struggle with the reasoning necessary for security analysis and detection consistently.

If not applied responsibly, generative AI can cause confusion by "hallucinating," where it references invented data, or by providing conflicting responses due to confirmation bias in the prompts written by different security team members.

---

# Darktrace machine learning: Combining multiple techniques

Each type of machine learning technique has its own set of strengths and weaknesses, so a multi-layered, multi-method approach is ideal to enhance functionality while overcoming the shortcomings of any one method. Darktrace's Cyber AI technology is powered by multiple machine learning approaches, which operate in combination for cyber defense.

**This allows Darktrace to protect the entire digital estates of the organizations it secures, including corporate networks, cloud computing services and SaaS, IoT, Industrial Control Systems (ICS), and email systems.**

## Multi-layered AI

Darktrace combines various machine learning types to create the AI that powers its products across the Darktrace ActiveAI Security Platform.

Plugged into the organization's infrastructure and services, our AI ingests and analyzes the data and its interactions within the environment and forms an understanding of the normal behavior of that environment, right down to the granular details of specific users and devices. The system continually revises its understanding about what is normal based on evolving evidence.

**This dynamic understanding of normal means that the AI engine can identify, with a high degree of precision, events or behaviors that are both anomalous and unlikely to be benign.**

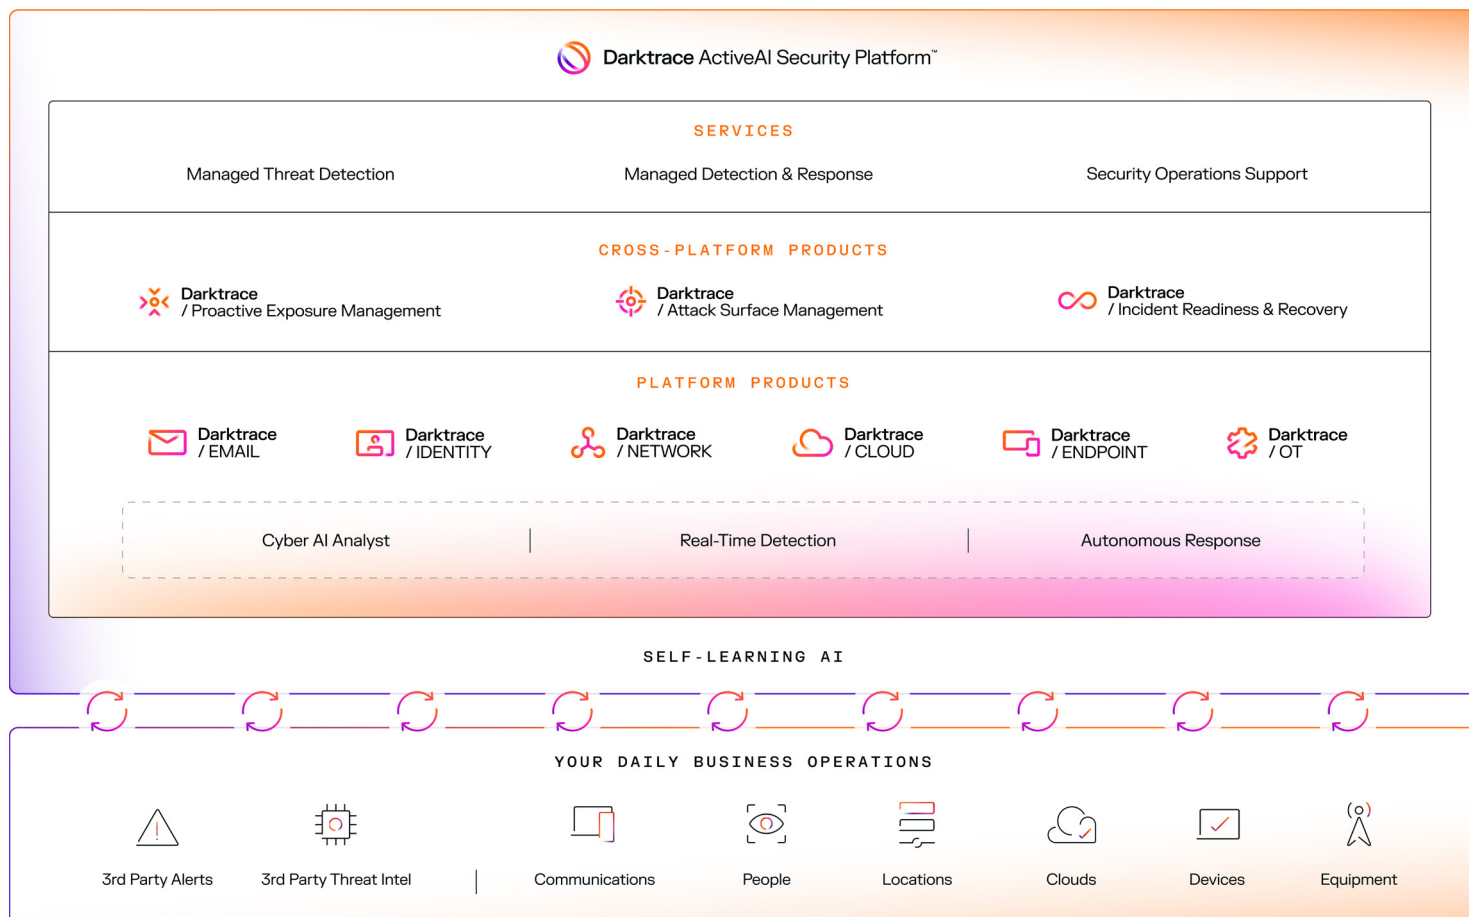


Figure 01: The Darktrace ActiveAI Security Platform delivers detection, response, and proactive security across traditionally siloed datasets

The ability to identify activity that represents the first footprints of an attacker, without any prior knowledge or intelligence, lies at the heart of the AI's efficacy in keeping pace with today's threat actors. It helps the human team detect subtle indicators that can be hard to spot amid the immense noise of legitimate, day-to-day digital interactions.

#### The often-unnoticed threats detected by Darktrace include:

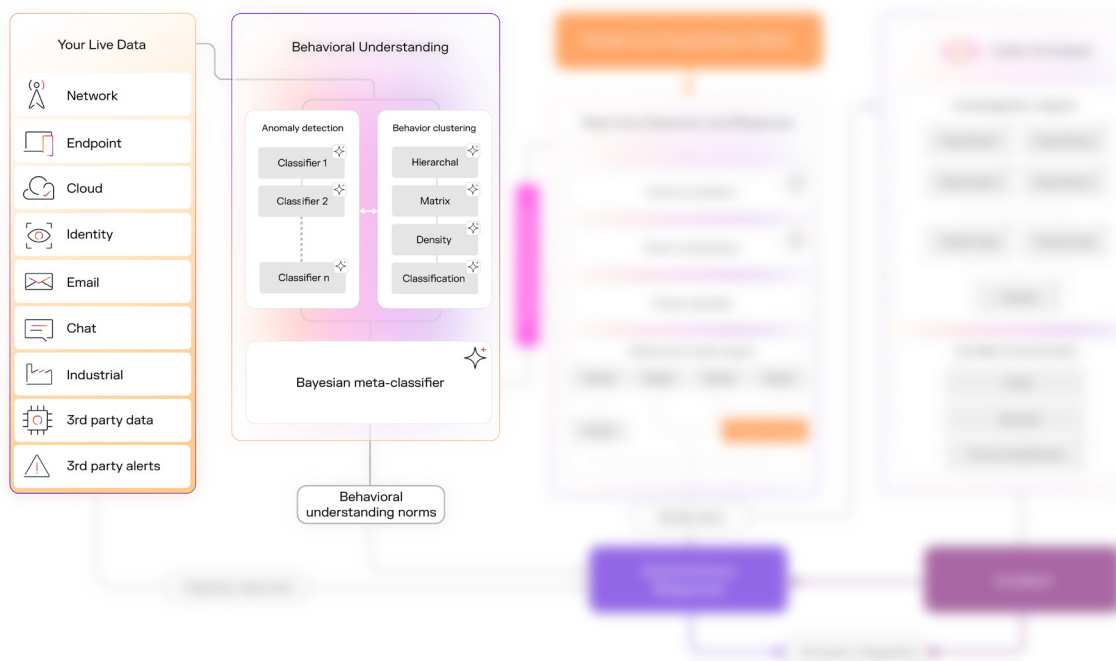
- [Insider threat](#) – malicious or accidental
- [Zero-day attacks](#) – previously unseen, novel exploits
- Latent vulnerabilities, whether [internal](#) or [external](#)
- Machine-speed attacks – ransomware and other automated attacks that propagate and/or mutate very quickly
- [Cloud and SaaS-based attacks](#)
- [Silent and stealthy attacks](#)
- [Advanced spear-phishing](#)

Let's take a closer look at how Darktrace uses AI to detect and respond to this wide range of threats. The following is a specific example of Darktrace's AI and how it combines supervised and unsupervised machine learning techniques as well as integrates the insights of our global deployments to improve threat detection.

Some platform products, such as Darktrace / EMAIL, / Attack Surface Management, and / Proactive Exposure Management, use slightly different AI architectures. However, the systems generally take the same AI approach, in that they use multiple kinds of AI to understand the unique environment.



# 1. Behavioral prediction



## Training data

Darktrace ingests live data from across your specific digital environment to continuously train and it integrates with third-party data and alerts to build a fuller picture. Notably, with this type of data, our security becomes tailored for each specific deployment. Other AI-powered tools learn on large data lakes to train general models, which can lead to oversimplifications and assumptions. By training on your organization's particular data, Darktrace AI is not comparing businesses and extrapolating behavioral differences between them. Instead, each AI system is unique to each individual business, creating high-fidelity detections.

## Bayesian probabilistic methods

To combine multiple analyses of digital activity, Darktrace uses Bayesian models that update with new information as it becomes available to the system or changes over time. Bayesian probabilistic methods involve unsupervised machine learning techniques that apply probabilistic modeling to both historical and current information. Constantly adjusting the probabilistic modeling creates a 'pattern of life' for assets, peer groups, and the organization as a whole.

Darktrace uses Bayesian probability to discover previously unknown relationships, independently classify data – even unstructured data without labels, and detect compelling patterns of what is normal and abnormal. This technique enables the continuous learning and the revising of assumptions about the behavior specific to the organization. Continually recalculating threat levels in light of new data, Darktrace can discern significant patterns in data flows indicative of attacks, where conventional signature-based methods see only chaos.

## Clustering algorithms

To model what should be considered as normal for an entity or asset, Darktrace analyzes its behavior in the context of other similar entities on the network. Darktrace uses unsupervised machine learning to algorithmically identify significant groupings, a task which is difficult and costly to do manually.

To create a holistic image of the relationships within the environment, Darktrace employs a range of different clustering methods including matrix-based, density-based, and hierarchical. The resulting clusters are then used to algorithmically identify significant groupings and inform modeling of normative behavior.

Continuous clustering not only identifies emerging suspicious activity before it looks obviously malicious, but it also identifies preexisting compromises. Other solutions that have definite training periods could learn malicious behavior as part of the baseline. However, clustering compares access, roles, identities, and functions across peer groups and so will recognize if no other devices are executing that same type of behavior.

“Within one week of installing Darktrace, it notified us to threats and vulnerabilities we had been totally unaware of.”

■ former CTO

Bunim/Murray Productions

---

## Anomaly detection models and Bayesian meta-classifier

The multi-layered AI accounts for ambiguities by distinguishing between the subtly differing levels of evidence that characterize operational data across the digital estate. Instead of generating the simple binary outputs of 'malicious' or 'benign,' these mathematical algorithms produce outputs marked with differing degrees of potential threat, including risk scores, rarity scores, and feature importance. This enables users of the system to rank alerts in a rigorous manner and prioritize those which most urgently require action.

---

At its core, Darktrace mathematically characterizes what constitutes 'normal' behavior, based on the analysis of many different measures of a device's behavior, including but not limited to:

- Protocol usage
- Frequency of connections
- Peer-to-peer communication
- Application usage patterns
- External communication patterns
- Port activity
- Traffic encryption
- Session duration
- Peripheral device usage
- DNS request patterns
- Email activity
- Interaction with suspicious entities
- Multi-factor authentication (MFA) usage

**Most anomaly detection models are only accurate if they are applied to the right data type. That means you cannot apply one anomaly detection machine learning model to all devices and all data types, because then you would have an accuracy problem.**

---

“As we shifted to the new mode of operation with people being remote, Darktrace very quickly gave us the ability to have the same functionality that we had when everybody was working on campus.”

■ **former CIO**

Salve Regina University

---

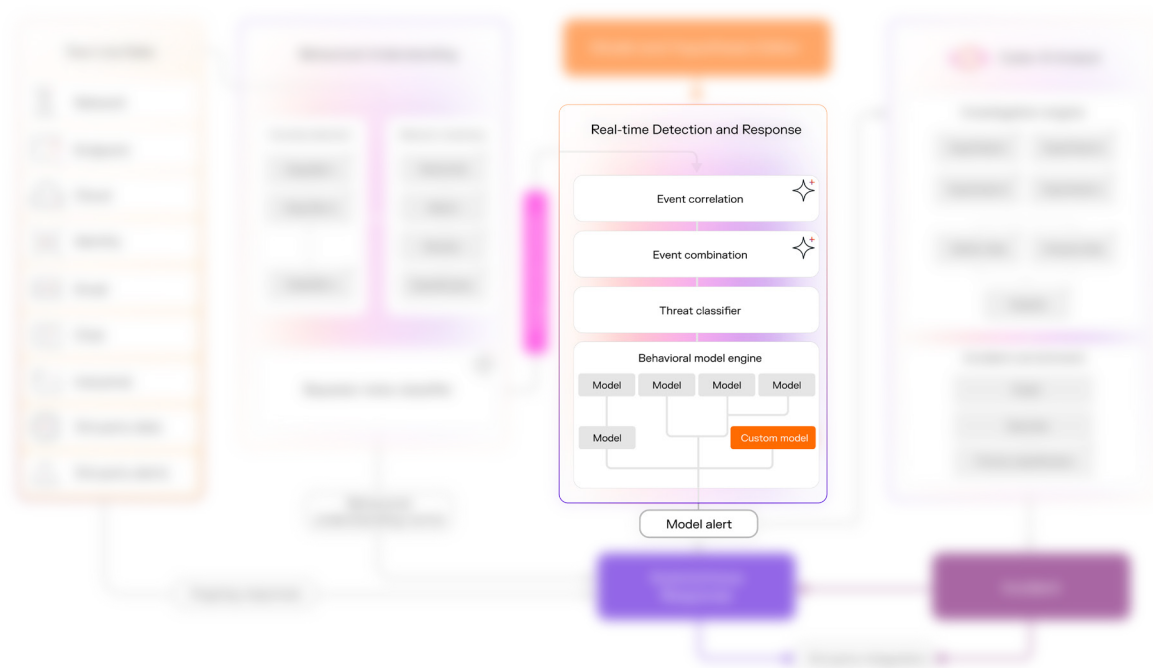
Given that, Darktrace takes an ensemble approach, scaling anomaly detection across hundreds of metrics. We then apply the Bayesian meta-classifier, which fine-tunes the underlying unsupervised machine learning models as well as provides the ongoing learning by making continual readjustments.

The Bayesian meta-classifier combines the outputs from multiple classifiers to produce an overall output and refine the results. For example, if a classifier is over firing, meaning it is not accurate for the environment or for the data it is ingesting, then the Bayesian meta-classifier will deprecate it and reduce the weight of importance given to its alerts for that specific deployment.

**As a result, some classifiers are more emphasized in each deployment, and that combination will be unique and tailored to the unique environment.**

To illustrate this, we can think of the analogy of journalism: more diversity of sources yields better corroboration and accuracy. By pairing the anomaly detection models with the Bayesian meta-classifier, the AI outputs become both collaborative and adaptive, yielding more accurate outputs.

## 2. Real-time threat detection



### Probabilistic and decision tree models

Event correlation and combination is performed to see broader scope of activity associated. With a true understanding of normal, our AI engine connects anomalous events to risky behavior. Anomalous behaviors are evaluated against the MITRE ATT&CK Framework to provide threat behavior context to anomalies.

### 3. Customizable model editor



#### Logic for processing control

On top of the machine learning models for detection, there is a layer of customizable logical statements, conditions, and models that perform threat classification and contextualization. This “Model Engine” layer gives the security team more visibility into the processing function of the output of the AI engine as well as the opportunity to customize outputs, therefore increasing explainability, interpretability, control, and the ability to modify the operationalization of the AI output with auditing. While the model editor comes with standard models (that are extensively updated), teams can edit this logic layer to control the AI engine which does the processing. The model editor allows teams to specify values, priorities, thresholds, and actions. That means a team can create custom detections based on use cases or business requirements. Teams can also create new detection models or increase the priority of existing detections to modify and control the AI’s behavior.

For example, some teams have used the model editor to customize default functions in order to profile a device, assign tags, highlight misconfigurations, or trigger autonomous responses. These custom models can integrate Darktrace into the existing security processes or replicate a SOC playbook.

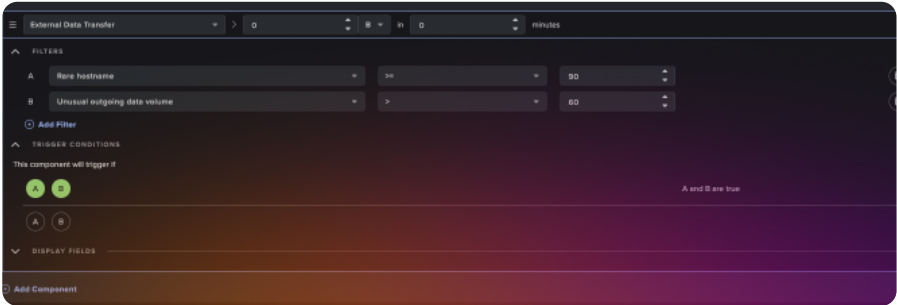


Figure 02: Alerts can be customized in the model editor in many ways like editing the thresholds for rarity and unusualness scores above.

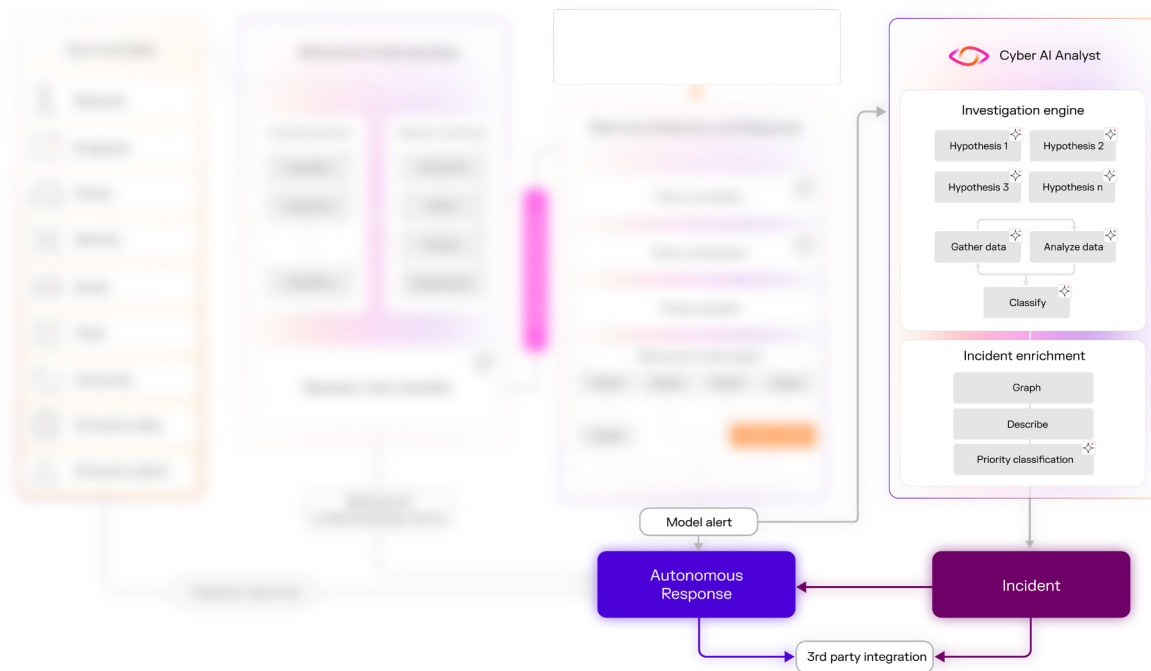
Notably, the model editor prevents model tampering, which is a risk to publicly hosted model gardens used by other AI vendors. This is because the models are not publicly hosted, they have attributed change control, and they have transparent, explainable logic so users can see whether changes have been made.

Even if teams do not customize models, the model editor ensures human users have visibility and understanding over the AI decision making. The various AI-powered capabilities show their work with clearly reported scores or explanations. For example, Darktrace specifies the threshold of what triggered a certain alert and the details of the AI’s investigations to see how it reached its conclusions.

**High interpretability of AI cybersecurity tools uplevels and upskills the human security team with more information to drive investigations and remediation actions. It also builds trust between the human security team and its AI tools, which is essential to maximize return on investment.**



## 4. Investigation and response



### Cyber AI Analyst

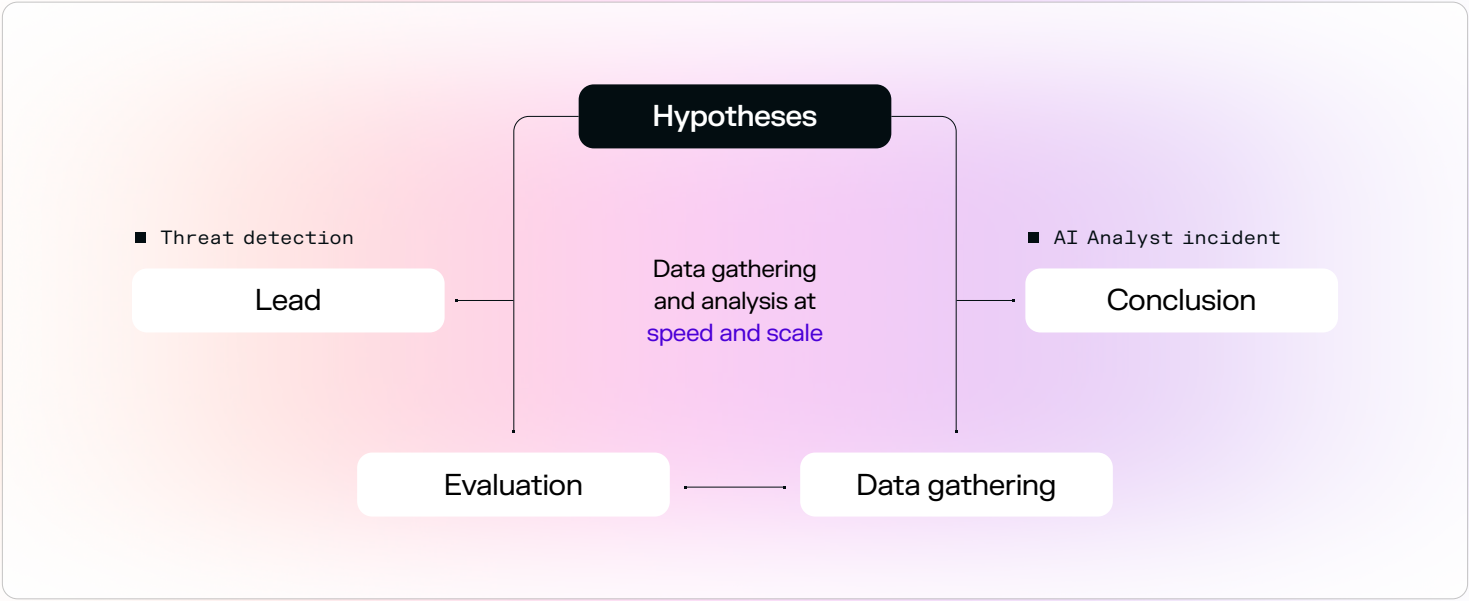
Darktrace uses various machine learning techniques to perform in-depth analysis and investigation of anomalies identified by models, in particular automating Levels 1 and 2 of a SOC team. This saves teams time and resources by helping to automate repetitive and time-consuming tasks carried out during investigation workflows. We call this core capability Cyber AI Analyst.

This investigative AI examines all relevant alerts. It uses unsupervised and supervised machine learning, with one source of its training data consisting of insights and behavior of our own expert Darktrace analysts.

**In this way, Cyber AI Analyst mimics the way a human carries out the threat investigation process. It**

- Investigates one or more hypotheses for relevant alerts
- Analyzes the logs of the assets involved
- Correlates incidents across different domains with graph-theoretic analysis
- Evaluates analysis against hypotheses – confirms, refutes, or modifies the hypothesis (and reiterates analysis as needed)
- Produces an alert with full technical details, which is part of a higher-level incident that tracks the whole lifecycle of a compromise
- Crafts an investigation summary including prioritization scores and an explanation of how the AI came to its conclusion

Cyber AI Analyst's process:



In the initial stages of the investigation process, the technology makes at least one broad hypothesis about what is happening, then queries and analyzes information as a human would – using custom algorithms and other machine learning techniques.

For example, it may use an LLM to analyze the hostname purpose, supervised machine learning to identify whether hostnames have been generated with a Domain Generation Algorithm (DGA), or unsupervised machine learning to compare against patterns of normal communications. The list goes on.

Cyber AI Analyst also uses graph theory to map relationships between assets within an organization's digital estate. This brings mathematical computations and weight based on importance, frequency, and similarities between assets.

**Unlike generative AI-based chatbot solutions, which apply semantic analysis to every kind of data, Cyber AI Analyst uses specific models to evaluate specific types of data. As a result, it can make intelligent decisions based on thousands of data points to save time and automatically investigate alerts.**

**Cyber AI Analyst typically provides SOC teams with the equivalent of up to 30 additional full time employees performing Level 2 analysis and written reporting annually, enriching operations by producing high level incident alerts with full details so that human analysts can focus on Level 3 tasks.**

Once the investigation has been concluded and Cyber AI Analyst understands the threat, the results can be classified using supervised machine learning to determine incidents of interest, at a speed and scale only possible with AI.

This allows for increasingly efficient and simplified investigations for analysts of all maturity levels. It also gives security teams the crucial time they need to focus on higher-value strategic work, such as managing risk and focusing on broader improvements to the business.

Cyber AI Analyst can often detect details that a human might miss, or might not have time to identify, and can decide whether or not an initial hypothesis holds in a matter of minutes. More crucially, the technology can classify and store the results of these investigations, allowing for only a small number of high priority incidents to be presented at any one time to reduce alert fatigue.

It communicates its findings and recommendations in the user interface, and additionally as detailed PDF reports, which present only a few high-priority incidents at any one time in natural language. It also produces alerts which can integrate with a wide range of different systems and services. These alerts and findings are then enriched with context and security insights that can be reviewed and understood by executives and end-users alike.

Importantly, Cyber AI Analyst adapts to new and unprecedented situations on the fly, enabling users to spend less time trawling through alerts, and more time prioritizing the strategic work that matters.

Autonomous Response

Darktrace's multi-layered AI does not only detect cyber-attacks, whether they are known or novel. It can also respond to them. Autonomous Response, a core capability of the Darktrace ActiveAI Security Platform, calculates the best action to take to neutralize in-progress attacks at machine speed.

As a result, cyber-attacks are contained until the security team can perform an investigation and remediation. It ensures that organizations are protected 24/7, even when the human team is out of office.

Autonomous Response directly interacts with the detections described previously. It can be triggered by model alerts and Cyber AI Analyst upon a significant deviation from the derived 'normal' for the device and its peer group, by the detection of specific malicious indicators or unwanted activities, or by a combination of small but meaningful indicators and subtle deviations from expected behavior.

This technology is supported by 'pattern of life' detection determined by unsupervised modeling and clustering, as well as a range of cutting-edge unsupervised classifiers that measure associations between users, activity patterns, and user intents.

Autonomous Response can generate proportionate, surgical responses to deliver these precise actions without disrupting daily business operations. Actions are targeted to the source of the threat and escalated only when necessary. For instance, Autonomous Response may sever an unusual File Transfer Protocol (FTP) connection or block access to Office 365 from an anomalous IP range. With its combination of varied AI techniques, this solution learns passively from the data that it observes as well as from itself. For example, when Autonomous Response generates an action, a feedback reinforcement loop is triggered. Used in this way, Cyber AI technology does not replace the human's function, but rather serves to enhance it. Autonomous Response acts faster than a human, buying the security team precious time to catch up.

“Darktrace provides us with protection and we can use it to make sure we’re as well-defended as we can be. In the future, organizations won’t be judged on suffering a cyber incident – because it’s almost inevitable. They’ll be judged on how they recover from incidents, and that’s where we know Darktrace, and the people that work there, have our backs.”

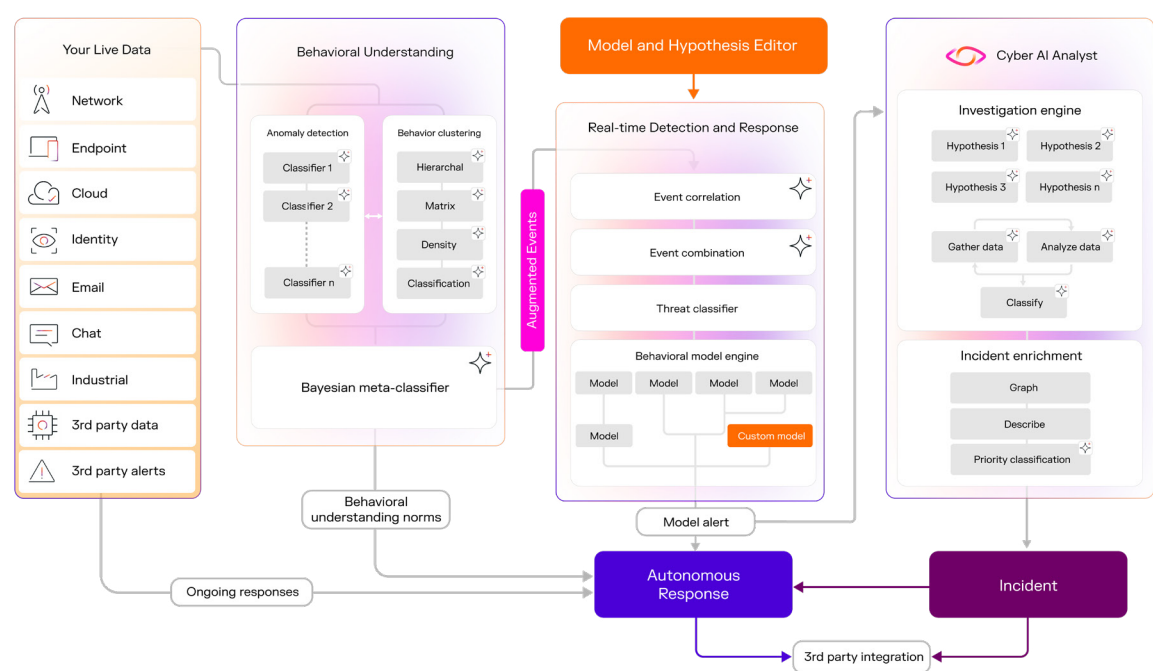
■ IT Manager  
Bristows

“We no longer live in an era where cyber-attacks are limited to the desktop or server. Darktrace’s machine learning fights the battle before it has begun.”

■ CIO  
City of Las Vegas

Fast and targeted action	Fully customizable	Protecting all digital environments
Responds to known or unknown attacks in seconds	<ul style="list-style-type: none"><li>Can operate fully autonomously, only with human confirmation, or a combination of the two</li><li>Integrates with firewalls and other security tools via API to take tailored actions</li></ul>	Applies across network, email, cloud, identities, endpoint, and Operational Technology (OT) as a core component of the Darktrace ActiveAI Security Platform

# Pulling the AI layers together



**Our multi-layered AI comes together to achieve behavioral prediction, real-time threat detection and response, and incident investigation, all while empowering your security team with visibility and control.**

We are in the Era of AI, experiencing major shifts in working practices with the widespread adoption of new technologies that can take on low value, repetitive tasks with automation and autonomous solutions capable of handling big data and making vast calculations. Security teams are challenged to keep up with a rapidly evolving cyber-threat landscape, now powered by AI in the hands of attackers, alongside the growing scope and complexity of digital infrastructure across the enterprise.


Traditional security methods, even those that use some simple machine learning, are no longer sufficient, as these tools can neither keep up with all possible attack vectors nor respond fast enough to the variety of machine-speed attacks, given their complexity compared to known and expected patterns. Security teams require a step up in their detection capabilities, using machine learning to understand the environment, filter the noise, and take action where threats are identified.

Darktrace's multi-layered AI technology has become a vital tool for security teams attempting to understand the scale of their networks, observe levels of activity, and detect areas of potential weakness. Machine learning technology is the fundamental ally in the defense of systems from the innovative hackers and insider threats of today, and in formulating responses to unknown methods of cyber-attack. It is a momentous change in cybersecurity, designed to empower security teams.

“Darktrace’s technology, experience, and expertise are helping us staying ahead of cyber-attacks, minimizing our risk and driving greater productivity for our team.”

■ CISO  
Severfield

**Learn more about how Darktrace uses AI across specific environments with these solution briefs:**

 **Darktrace / NETWORK**

 **Darktrace / CLOUD**

 **Darktrace / EMAIL**

[Contact](#) Darktrace for a personalized demo



■ About Darktrace

Darktrace is a global leader in AI cybersecurity that keeps organizations ahead of the changing threat landscape every day. Founded in 2013 in Cambridge, UK, Darktrace provides the essential cybersecurity platform to protect organizations from unknown threats using AI that learns from each business in real-time. Darktrace's platform and services are supported by 2,400+ employees who protect nearly 10,000 customers globally. To learn more, visit <http://www.darktrace.com>.